NASA/CR-1999-209826
ICASE Report No. 99-49

Optimization with Variable-fidelity Models Applied to Wing Design

Natalia M. Alexandrov
NASA Langley Research Center, Hampton, Virginia

Robert Michael Lewis
ICASE, Hampton, Virginia

Clyde R. Gumbert, Larry L. Green, and Perry A. Newman
NASA Langley Research Center, Hampton, Virginia

Institute for Computer Applications in Science and Engineering
NASA Langley Research Center
Hampton, VA
Operated by Universities Space Research Association

National Aeronautics and
Space Administration

Langley Research Center
Hampton, Virginia 23681-2199

Prepared for Langley Research Center
under Contract NAS1-97046

December 1999

OPTIMIZATION WITH VARIABLE-FIDELITY MODELS APPLIED TO WING DESIGN

NATALIA M. ALEXANDROV*, ROBERT MICHAEL LEWIS†, CLYDE R. GUMBERT‡, LARRY L. GREEN§, AND PERRY A. NEWMAN¶

Abstract. This work discusses an approach, the Approximation Management Framework (AMF), for solving optimization problems that involve computationally expensive simulations. AMF aims to maximize the use of lower-fidelity, cheaper models in iterative procedures with occasional, but systematic, recourse to higher-fidelity, more expensive models for monitoring the progress of the algorithm. The method is globally convergent to a solution of the original, high-fidelity problem. Three versions of AMF, based on three nonlinear programming algorithms, are demonstrated on a 3D aerodynamic wing optimization problem and a 2D airfoil optimization problem. In both cases Euler analysis solved on meshes of various refinement provides a suite of variable-fidelity models. Preliminary results indicate threefold savings in terms of high-fidelity analyses in case of the 3D problem and twofold savings for the 2D problem.

Key words. approximation concepts, approximation management, model management, surrogate optimization, aerodynamic optimization, nonlinear programming, wing design

Subject classification. Applied and Numerical Mathematics

1. Introduction. Many physical phenomena in engineering design can be described by computational models of high physical fidelity or numerical accuracy. However, the use of high-fidelity models, such as the Navier-Stokes equations or those based on fine computational meshes, in iterative procedures can be prohibitively expensive. On the other hand, the use of corresponding lower-fidelity models alone does not guarantee improvement for higher-fidelity design. This paper discusses an approach that aims to facilitate design optimization and integration of disciplines in a multidisciplinary environment by alleviating the expense of relying exclusively on high-fidelity models, while taking advantage of well-established engineering approximation concepts.

Computational models of varying fidelity have been used in engineering design for a long time [32, 33, 34]. A survey on the use of approximations in structural optimization can be found in Barthelemy and Haftka [9]. Recent overviews of algorithms for aerodynamic analysis and optimization can be found, e.g., in Jameson [23] and Newman et al. [28]. Procedures for using variable-fidelity models, however, have been largely based on heuristics, and convergence to a solution of the highest-fidelity optimal design problem has not been guaranteed, in general. Until recently, with a few exceptions [8, 27, 29], the analysis had focused on the question of convergence to a solution of the surrogate problem [14, 22]. Due to improvements in the numerical

*Multidisciplinary Optimization Branch, Mail Stop 159, NASA Langley Research Center, Hampton, VA, 23681-2199, n.alexandrov@larc.nasa.gov.

†ICASE, Mail Stop 132C, NASA Langley Research Center, Hampton, Virginia 23681-2199, buckaroo@icase.edu. This research was supported by the National Aeronautics and Space Administration under NASA Contract No. NAS1-97046 while the author was in residence at the Institute for Computer Applications in Science and Engineering (ICASE), NASA Langley Research Center, Hampton, VA 23681-2199.

‡Multidisciplinary Optimization Branch, Mail Stop 159, NASA Langley Research Center, Hampton, VA, 23681-2199, c.r.gumbert@larc.nasa.gov.

§Multidisciplinary Optimization Branch, Mail Stop 159, NASA Langley Research Center, Hampton, VA, 23681-2199, l.l.green@larc.nasa.gov.

¶Multidisciplinary Optimization Branch, Mail Stop 159, NASA Langley Research Center, Hampton, VA, 23681-2199, p.a.newman@larc.nasa.gov.

modeling techniques and the increased availability of high-fidelity analyses, optimization with variable-fidelity approximations has become a subject of much interest in the past few years [15, 21, 30].

A number of methods for managing models and approximations of varying fidelity have been recently introduced and analyzed. These include methods that use sensitivities [3, 25, 1, 5] and methods that do not [13]. This paper begins a computational investigation of the practical effectiveness of the methods in [3, 25, 1, 5] on problems of aerodynamic design optimization.

For the purposes of the present work, the optimal design problem is represented by the following nonlinear programming problem (NLP):

$$
\begin{array}{ll}
\text{minimize} & f(x) \\
\text{subject to} & g(x) \le 0 \\
 & l \le x \le u,
\end{array}
\tag{1.1}
$$

where x are the design variables, the objective f and the vector-valued constraints g are smooth (i.e., continuously differentiable) nonlinear functions, and $l \le x \le u$ denotes bound constraints on design variables. This work concerns a general approach for controlling the use of variable-fidelity models—the first-order trust-region Approximation Management Framework (AMF)—in solving problem (1.1).

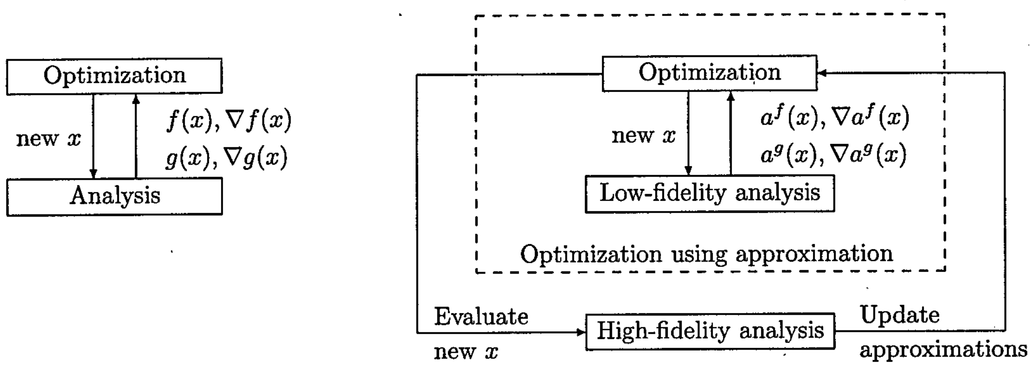

FIG. 1.1. *Conventional optimization vs. AMF*

The conceptual distinction between AMF and conventional optimization is depicted in Figure 1.1. On the left, in conventional optimization, the optimizer and the analysis software exchange information as follows. The analysis supplies the optimizer with objective and constraint function and derivative information, $f, \nabla f, g, \nabla g$, while the optimizer produces new values of the design variables x for re-analysis. The optimizer uses the function and derivative information to build local approximations—usually first or second-order Taylor series—internally. If evaluating the problem functions and derivatives involves a simulation of high accuracy but high computational cost (e.g., the Navier-Stokes equations), repeated consultations with analysis required by the optimizer are expensive.

Now suppose one also has a suite of less accurate but also less expensive approximate models or surrogates[1] of the same phenomenon. A lower-fidelity model of the objective is denoted by $\{a^f(x)\}$ and lower-

[1]Some authors make distinctions in the use of the terms "models", "surrogates", and "approximations." For simplicity, the terms are used interchangeably here.

fidelity model of the constraints by $\{a^g(x)\}$. The associated sensitivities with respect to the design variables are $\{\nabla a^f(x)\}$ and $\{\nabla a^g(x)\}$.

The right-hand side of Figure 1.1 depicts information exchange between the optimizer and the analyses in a conceptual AMF scheme. Here the optimizer receives the function and sensitivity information $a^f(x), \nabla a^f(x), a^g(x), \nabla a^g(x)$ from the lower-fidelity model to build internal local models (Taylor series). Expensive, high-fidelity computations proceed outside the optimization loop and serve to re-calibrate the lower-fidelity model occasionally, based on a set of systematic criteria. The salient features of AMF are as follows.

- Although a lower-fidelity model may not capture a particular feature of the physical phenomenon to the same degree of accuracy (or at all) as its higher-fidelity counterpart, a lower-fidelity model may still have satisfactory predictive properties for the purposes of finding a good direction of improvement for the higher-fidelity model.
- AMF replaces the local, Taylor series approximations of the conventional optimization by general nonlinear models required to satisfy a set of first-order consistency conditions defined later in the paper. In principle, AMF is capable of handling arbitrary models, provided the easily imposed consistency conditions are satisfied.

  In particular, AMF is not limited to the use of algebraic, Taylor series, or response surface approximations. Analyses of variable mesh resolution or of variable physical fidelity (such as Navier-Stokes and Euler codes) can be used as variable-fidelity function evaluators in AMF.
- In AMF-based algorithms, the bulk of the computational expense involves calculations based on lower-fidelity models in iterations of optimization or search procedures.
- AMF is based on the trust-region methodology [18, 26], which can be described as an adaptive move limit strategy for improving the global behavior of optimization algorithms based on local models. The trust-region methodology ensures the convergence of the AMF scheme to a solution of the higher-fidelity problem by providing a measure of the surrogate's predictive behavior, a criterion for updating the surrogate, and a systematic response to situations in which an optimization phase performed using a surrogate gives either an incorrect or a poor prediction of the higher-fidelity model's actual behavior.

The proposed AMF methods have strong theoretical properties and have been tested on small, closed-form problems. The framework's generality means, however, that it admits not only a wide range of models, but also an extensive set of rules for governing parts of AMF. The rules do not influence the fact of convergence, but they strongly influence the algorithm's efficiency. Moreover, the performance of the algorithm will also be problem-dependent. Hence, much computational experience is needed both to validate the concept and to arrive at an advantageous set of rules for the use of approximations within a specific instance of AMF. This study intends to demonstrate the concept of AMF on two aerodynamic optimization problems. In addition, the study begins the accumulation of the necessary computational experience that should result in further practical implementations of AMF, both for single-discipline and multidisciplinary design optimization (MDO).

The paper is organized as follows. The next section briefly describes the three AMF under study. The computational demonstration is then described. The paper concludes with lessons learned to date and some mention of ongoing work.

2. AMF Under Study. There are, in principle, as many possible AMF as there are optimization algorithms, because AMF may be defined as a mechanism for a systematic alternation between the available

models within an optimization procedure. A detailed discussion of the algorithms and analysis of the entire class of first-order AMF for constrained optimization can be found in Alexandrov and Lewis [7].

In this section, three AMF are described, and the reasons for their selection are discussed. The first AMF is based on an augmented Lagrangian approach, the second on a multi-level optimization algorithm, and the third on an SQP approach. In the remainder of the paper, the subscripts "c" and "+" denote the current and the new iterate, respectively.

2.1. Augmented Lagrangian-based AMF. The augmented Lagrangian method for constrained optimization allows for an immediate extension of the unconstrained AMF to constrained problems. The underlying algorithm is the augmented Lagrangian approach as implemented in Conn, Gould, and Toint [17].

In this method, the explicit nonlinear inequality constraints of problem (1.1) are converted to equalities by introducing nonnegative artificial (slack) variables z to define the equality constraints

$$h(x,z) \equiv g(x) + z,$$

which gives rise to the following equivalent formulation:

$$\begin{array}{ll} \underset{x,z}{\text{minimize}} & f(x) \\ \text{subject to} & h(x,z) = 0 \\ & l \leq x \leq u \\ & z \geq 0. \end{array} \tag{2.1}$$

The associated augmented Lagrangian is

$$L(x,z,\lambda;\mu) \equiv f(x) + \lambda^T h(x,z) + \frac{1}{2\mu} \parallel h(x,z) \parallel_2^2,$$

where λ is the vector of Lagrange multipliers, and $\mu > 0$ is the penalty parameter. For appropriate values of μ and λ, minimization of L solves problem (2.1). However, since the appropriate values of μ and λ are not known *a priori*, an iterative approach is devised that solves an augmented Lagrangian subproblem while updating μ and λ. The conventional augmented Lagrangian approach is described in the following pseudo-code:

Initialize (x_c, z_c), λ_c, μ_c
Do until convergence:
With (x_c, z_c) as the initial point and fixed λ_c, μ_c,
solve the following subproblem for (x_+, z_+):

$$\begin{array}{ll} \underset{x,z}{\text{minimize}} & L(x,z;\lambda_c,\mu_c) \\ \text{subject to} & l \leq x \leq u \\ & z \geq 0 \end{array}$$

Set $(x_c, z_c) = (x_+, z_+)$
Update λ_c and μ_c
End do

The corresponding augmented Lagrangian-based AMF uses a subproblem that minimizes not L but its model a_c^L, which yields the following AMF:

Initialize (x_c, z_c), λ_c, μ_c
Do until convergence:
Compute the high-fidelity L and ∇L at (x_c, z_c)

Select an approximation a_c^L to L, with
$a_c^L(x_c, z_c, \lambda_c; \mu_c) = L(x_c, z_c, \lambda_c; \mu_c)$ and
$\nabla a_c^L(x_c, z_c, \lambda_c; \mu_c) = \nabla L(x_c, z_c, \lambda_c; \mu_c)$
Do until convergence:
Solve approximately for (x_+, z_+):

$$\begin{array}{ll} \underset{x,z}{\text{minimize}} & a_c^L(x, z; \lambda_c, \mu_c) \\ \text{subject to} & l \leq x \leq u \\ & z \geq 0 \\ & \| x_c - x \|_\infty \leq \Delta_c \\ & \| z_c - z \|_\infty \leq \Delta_c \end{array}$$

End do
Compute $R = \dfrac{L(x_c, z_c, \lambda_c; \mu_c) - L(x_+, z_+, \lambda_c; \mu_c)}{L(x_c, z_c, \lambda_c; \mu_c) - a_c^L(x_+, z_+, \lambda_c; \mu_c)}$
Update Δ_c and (x_c, z_c), based on the value of R
Update λ_c and μ_c
End do

Minimizing a_c^L is itself an iterative procedure that now uses approximations (e.g., Taylor series) to the lower-fidelity model a_c^L. See reference [7] for further details.

The conditions on the model a_c^L are known as the first-order consistency conditions. They are imposed by a correction technique introduced by Chang et al. [16]. This technique corrects a low-fidelity version ϕ_{lo} of a function so that it agrees to first-order with a given high-fidelity version ϕ_{hi}. This is done by defining the scale factor β

$$\beta(x) = \frac{\phi_{\text{hi}}(x)}{\phi_{\text{lo}}(x)}.$$

Given the current design variables x_c, one builds a first-order model β_c of β about x_c:

$$\beta_c(x) = \beta(x_c) + \nabla\beta(x_c)^T(x - x_c).$$

The local model of β is then used to scale ϕ_{lo} to obtain a better approximation $a(x)$ of ϕ_{hi}:

$$\phi_{\text{hi}}(x) = \beta(x)\phi_{\text{lo}}(x) \approx a(x) \equiv \beta_c(x)\phi_{\text{lo}}(x).$$

The corrected approximation $a(x)$ has the properties that $a(x_c) = \phi_{\text{hi}}(x_c)$ and $\nabla a(x_c) = \nabla\phi_{\text{hi}}(x_c)$.

The augmented Lagrangian-based AMF is relatively easy to implement and can be proven to work reliably under reasonable assumptions. The underlying framework is well understood and is a basis for a number of popular codes. The expected difficulties are also those of the underlying optimization algorithm—augmented Lagrangian methods can converge slowly, depending on how λ is estimated, and they are also subject to ill-conditioning as μ approaches 0.

2.2. MAESTRO-based AMF. Another AMF is based on a class of multilevel algorithms for large-scale constrained trust-region optimization (MAESTRO) [2, 4, 6]. This AMF is of interest due to MAESTRO's convergence properties and a natural structure for MDO problems with arbitrary couplings. The reader is referred to Alexandrov [5] for details of the MAESTRO-based AMF. A brief description follows.

The present version of the underlying MAESTRO approach deals with problem (1.1) by converting the

explicit inequalities into equalities via squared slacks:

$$\begin{array}{ll} \underset{x,z}{\text{minimize}} & f(x) \\ \text{subject to} & h(x,z) \equiv g(x) + z^2 = 0. \end{array} \tag{2.2}$$

Because the current demonstration problems have one discipline and a small number of variables, only the bilevel MAESTRO procedure will be described here.

If (x_c, z_c) is the current iterate, a model of the constraints a_c^h is first selected that satisfies the following consistency conditions for the constraints at that point:

$$\begin{aligned} a_c^h(x_c, z_c) &= h(x_c, z_c) \\ \nabla a_c^h(x_c, z_c) &= \nabla h(x_c, z_c). \end{aligned} \tag{2.3}$$

A substep $s_1 = (s_1^x, s_1^z)$ is computed that approximately minimizes that model within a trust region. The process of computing the substep is itself an iterative procedure. Next, a model a_c^f of the objective function or the Lagrangian is selected that satisfies the consistency conditions at the just computed point:

$$\begin{aligned} a_c^f(x_c + s_1^x) &= f(x_c + s_1^x) \\ \nabla a_c^f(x_c + s_1^x) &= \nabla f(x_c + s_1^x). \end{aligned} \tag{2.4}$$

The substep $s_2 = (s_2^x, s_2^z)$ is computed in another loop to approximately minimize the model in another trust region. The total trial step s_c is the sum of the two substeps. The step is evaluated using a merit function (the augmented Lagrangian or the ℓ_2 penalty function).

The consistency conditions can be relaxed, but that line of reasoning will not be pursued here, because, given any two models, the conditions are easily enforced by using the correction techniques due to Chang, et al. [16]. The consistency conditions are not enforced within each optimization sweep.

This algorithm for computing the trial step is a special case of the MAESTRO class with the distinction that the Gauss-Newton model of the constraints and the quadratic model of the objective or the Lagrangian have been replaced by general, first-order models that satisfy the consistency conditions (2.3) and (2.4).

Since the underlying algorithm belongs to the MAESTRO class, this AMF will converge to a critical point of the high-fidelity problem under the assumptions that lead to convergence of the underlying class. This means finding the substeps s_1 and s_2 that will satisfy the sufficient decrease conditions necessary for establishing convergence.

Implementing a MAESTRO-based AMF is more laborious than the augmented Lagrangian-based AMF. The benefits are greater efficiency and the expected incorporation of MDO problems in the near future.

2.3. SQP-based AMF. Sequential Quadratic Programming (SQP) methods are a popular class of methods for solving nonlinear programming problems. An overview of these algorithms can be found in Gill, Murray, and Wright [20]. There are many variants of SQP. If B_c denotes an approximation to the Hessian of the objective function f, one conventional approach to solving problem (1.1) is the following:

Initialize x_c

Do until convergence:

Solve the following subproblem for $s_c = x - x_c$:

$$\begin{array}{ll} \underset{s}{\text{minimize}} & \nabla f(x_c + s)^T s + \frac{1}{2} s^T B_c s \\ \text{subject to} & g(x_c) + \nabla g(x_c)^T s \leq 0 \\ & l \leq x \leq u \end{array}$$

Update x_c

End do

Typically, the objective function in the SQP subproblem is a quadratic approximation to the problem Lagrangian. Globalization strategies, such as line search or trust-region approaches, are then used to insure the robustness of the algorithm.

An SQP-based AMF studied here is based on the algorithm just outlined. It has been selected for the study as an alternative to MAESTRO in the case of single-discipline optimization, where the objective and constraint values are given by a single analysis execution (i.e., when there is no natural multilevel structure in the problem).

Let $P(x;\mu)$ be a merit function for the high-fidelity problem. In the work described here, P is the l_1 penalty function, but other choices are possible. The SQP-based AMF is:

Initialize x_c, μ_c

Do until convergence:

Select approximations a_c^f and a_c^g, with

$a_c^f(x_c) = f(x_c)$; $\nabla a_c^f(x_c) = \nabla f(x_c)$ and

$a_c^g(x_c) = g(x_c)$; $\nabla a_c^g(x_c) = \nabla g(x_c)$.

Solve approximately for $s = x - x_c$:

$$\begin{array}{ll} \underset{s}{\text{minimize}} & a_c^f(x_c + s) \\ \text{subject to} & a_c^g(x_c) + \nabla a_c^g(x_c)^T s \leq 0 \\ & l \leq x \leq u \\ & \| s \|_\infty \leq \Delta_c \end{array}$$

End do

Compute $P(x_c + s_c)$

Update Δ_c and x_c, μ_c based on the value of P

End do

Details and analysis of the implementation can be found in [7]. Briefly, the approach has a number of benefits. The SQP-based AMF is relatively easy to implement and converges very rapidly once it is near a solution. It handles the inequality constraints directly and enjoys the efficiency of SQP methods. By choosing Δ_c sufficiently large, it can be arranged for the first iteration to go to a solution of the lower-fidelity problem. This feature must be obtained by pre-processing in the other approaches. The SQP-based AMF also allows for an easy incorporation of commercial software. The drawbacks of the approach are not obvious at this point.

3. Computational Demonstrations. The computational demonstrations are intended to validate the effectiveness of AMF. The ability to transfer the computational load onto the lower-fidelity, cheaper computations, and thereby reduce the overall computational cost, will depend on the predictive qualities of the surrogates. Note that even though the surrogate models may not be good approximators of the higher-fidelity models for the purposes of analysis, they may possess suitable predictive properties for the purposes of optimization. That is, an approximation may not capture all the important properties of a higher-fidelity function, but it may still produce a step that will lead to a satisfactory improvement (decrease or increase) in the merit function for the higher-fidelity problem.

The computational tests include both the case when the relationship between the various levels of models is favorable and the case when it is not. The relationship is favorable when the lower-fidelity model can provide a long sequence of steps with satisfactory directions of descent for the higher-fidelity merit function

before the lower-fidelity model has to be re-calibrated. The relationship is not favorable when the lower-fidelity model does not satisfactorily capture the trends in the objective and constraints computed using the higher-fidelity model on a significant region of the feasible region.

The AMF approaches could suffer from an over-reliance on the low-fidelity models if the lower-fidelity surrogate does not predict the behavior of the higher fidelity model adequately. In this case, the AMF will be forced to take only a few steps using the surrogate information before having to resort to re-calibrating the model, which, in effect, means optimization with high-fidelity models. Thus, in the worst case, the AMF approach reverts to conventional optimization with the high-fidelity models.

3.1. Computational Experiments. The tests described in this paper investigate a specific type of variable-fidelity modeling—that in which performing a single type of analysis (aerodynamic analysis using the Euler equations) on a variety of related meshes provides variable-fidelity models. In this case, the finer the mesh, the higher the model fidelity and computational expense.

It will be significantly more difficult to determine the suitability of the AMF concept when the variable-fidelity models are represented by different physical models, such as Navier-Stokes versus Euler, and particularly in extreme cases of varying the model fidelity, for instance, Navier-Stokes versus a linear panel code. It is likely that the framework will prove infeasible for some model combinations. The accumulation of numerical experience in the simpler case of variable-resolution models based on a single physical model will provide a necessary foundation for a study of the more difficult case of different physical models.

The initial experiments are conducted only for two design variables in order to visualize the algorithms' progress easily and completely.

The problems were first solved in a single-fidelity mode by using well-known commercial optimization software[2], such as NPSOL [19] and PORT [24], in order to obtain a baseline number of function evaluations or iterations to find an optimum. The problems were then solved in a single-fidelity mode with research implementations of methods on which AMF are based to obtain a baseline for comparison with AMF. The problems were finally solved with the AMF.

In the study of the two demonstration problems, a number of interesting issues concerning the quality and selection of models arose. Initially, computations were done on arbitrary meshes of different sizes with no relation between the meshes. While sufficiently fine meshes should, in principle, produce consistent functions, the meshes were too coarse (even the finest one) to observe this effect. Instead, objectives and constraints computed on unrelated meshes could have wildly disparate trends and features, a phenomenon observed by other investigators [12]. This difficulty was remedied by using coarser meshes that were proper subsets of the finest mesh.

Initial tests were conducted using the actual function evaluations obtained by executing the analysis software. An examination of the problem functions revealed that they exhibited benign behavior, insofar as the objectives and constraints in the present study are smooth and very nearly convex. However, function and gradient evaluation is very expensive even for the small number of design variables under consideration here. Because the underlying problem was benign but expensive to compute, it was decided to accumulate generous amounts of data and to replace the actual functions with a number of accurate response surfaces. It should be emphasized that the use of response surfaces is *not* an integral part of the approach, and is not even one of the focuses of this study. Response surfaces were introduced strictly to facilitate the testing, because they approximate the trends of the actual test functions so well at a tiny fraction of the computational cost.

[2]The use of names of commercial software in this paper is for accurate reporting and does not constitute an official endorsement, either expressed or implied, of such products by the National Aeronautics and Space Administration or ICASE.

An additional benefit of using response surfaces later became apparent. In particular, graphics will show that for the problems under study, the lower-fidelity functions obtained on coarser meshes provide an excellent approximation (with respect to descent characteristics) to those computed on finer meshes. This is a most favorable condition for AMF. However, one must also investigate cases where the lower-fidelity problem does not capture the high-fidelity descent behavior well. Some of the response surfaces provide such a test case.

In the experiments, surfaces based on data computed by executing analyses on finer meshes represents high-fidelity functions, while those based on data obtained from analyses on coarser meshes serve as low-fidelity functions. Again, the response surfaces are used in the experiments solely to reduce the computational cost to a point where testing and debugging are practical.

Three types of response surfaces were used:

1. Two-dimensional, uniform, variation diminishing splines (obtained from the PORT [24] package);
2. Kriging (implemented locally);
3. Cubic polynomial response surfaces (coded locally with assistance from the RSG [38] package).

For all three AMF, the consistency conditions were enforced via the scaling technique in Chang et al. [16]. This technique was found to provide an excellent correction strategy for the lower-fidelity model in all cases.

Performance of AMF's is evaluated in terms of the absolute number of calls to the high and low-fidelity function and sensitivity calculations and the number of "equivalent" high-fidelity computations. The latter are easily obtained because both analysis codes use multigrid techniques, where this metric is commonly computed.

Finally, a conscious effort was made to implement the AMF in a straightforward manner, without any "fine-tuning", in order to obtain a proof of concept. As will be discussed later, significant improvements in efficiency can likely be made.

4. 3D Wing Problem.

4.1. Optimization Problem. The first demonstration problem is a three-dimensional aerodynamic wing optimization. The wing consists of a single trapezoidal panel with a rounded tip. It is parameterized by fifteen variables, five of which describe the planform, five of which describe the root section shape, and five of which describe the tip section shape. The wing and some of the associated parameters are depicted in Figure 4.1. Currently, the two design variables are the tip chord and the tip trailing-edge setback. The objective function $f(x)$ is the negative lift-to-drag coefficient ratio, $-C_L/C_D$. Several artificial constraints are imposed in lieu of multidisciplinary constraints. Purely geometric constraints ensure a minimum leading edge radius and a minimum thickness. "Aerodynamic" constraints are:

1. A lower bound on total lift $C_L \times S$, in lieu of a minimum payload requirement (S is the semispan wing planform area);
2. An upper bound on C_M (pitching moment coefficient), in lieu of a trim constraint;
3. An upper bound on C_l (rolling moment coefficient), in lieu of a maximum bending moment.

For the results we report here, the flow is subsonic, with a free-stream Mach number $M_\infty = 0.5$. The angle of attack is $\alpha = 3°$. Given the subsonic speed of the flow, the drag is primarily the induced drag due to lift, $C_D \propto C_L^2$, so our objective is effectively $-1/C_L$.

The aerodynamic analysis code used for this study is CFL3D.ADII [37], a version of CFL3D [31] obtained via the ADIFOR automatic differentiation tool [10]. The surface geometry was computed based on the problem parameters via software that uses the RAPID technique [35]. The ADIFOR generated analysis

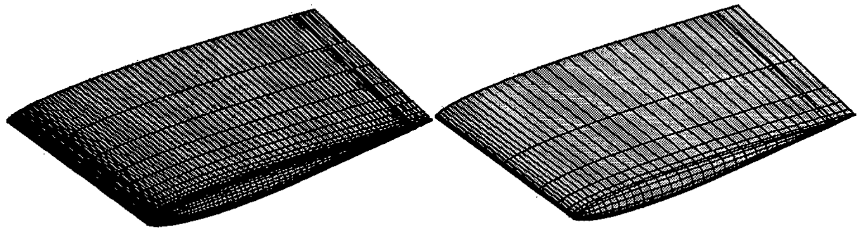

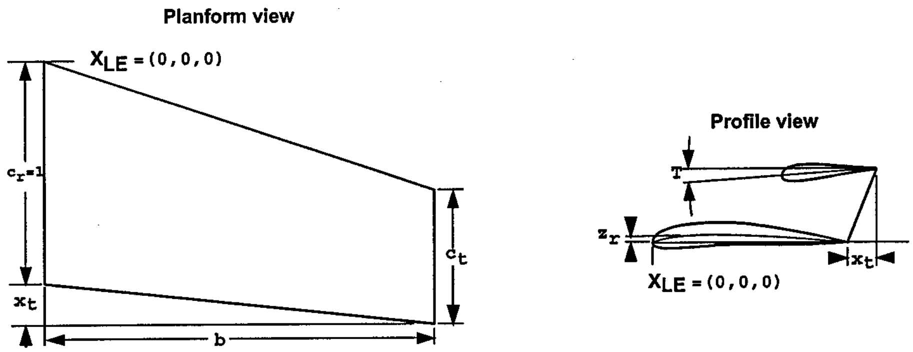

FIG. 4.1. *The 3D wing problem*

code includes the capability for computing the gradients. The volume mesh and associated gradients needed for CFL3D are generated using a ADIC [11] generated version of CSCMDO.

Two variable-fidelity models and associated constraints are generated by performing the CFL3D.ADII analysis on meshes of increased refinement:

1. 97×25×17 (low fidelity) and
2. 193×49×33 (high fidelity).

Since the analysis uses a multigrid solution process, the CPU time per converged function evaluation is essentially linear in the number of grid points, resulting in an eight-fold difference in execution time between adjacent levels of fidelity. For example, on an Ultra 1 Sun workstation, a single function and constraint evaluation on the 97×25×17 mesh takes eight minutes, and the 193×49×33 mesh analysis takes about an hour, without computing derivatives.

4.2. Discussion of Numerical Results. Figure 4.2 depicts the level sets of the objective functions and active constraints obtained by performing analyses on the 193×49×33 and 97×25×17 meshes. The shaded regions are infeasible. For the current, subsonic case, the rolling moment constraint C_l is inactive and is not depicted. Solutions are marked by black squares. Note that this problem has a favorable structure for AMF. Although the optima are at different locations, the low-fidelity and high-fidelity objective and constraints have similar trends.

Initial testing on this problem was done with MAESTRO-based AMF and with function values obtained directly from CFL3D.ADII on the 193×49×33 mesh for high-fidelity and 97×25×17 for low fidelity. For the case studied, none of the constraints were active. The analysis count was as follows. To obtain a solution on the low-fidelity mesh alone, using non-AMF MAESTRO, required 17 function and 17 sensitivity calls. Solution with the high-fidelity mesh alone was attempted but not completed, due to the expense of direct function and derivative evaluations. However, it is reasonable to assume that the solution would not have taken fewer iterations than that on the low-fidelity mesh. The MAESTRO-based AMF required 18 low-

fidelity functions, 18 low-fidelity sensitivities, 7 high-fidelity functions, and 7 high-fidelity sensitivities, for a total of $7 + 18/8 = 9\ 1/4$ equivalent high-fidelity functions and as many sensitivities. Thus the increase in efficiency is approximately twofold, both in the number of function and sensitivity computations.

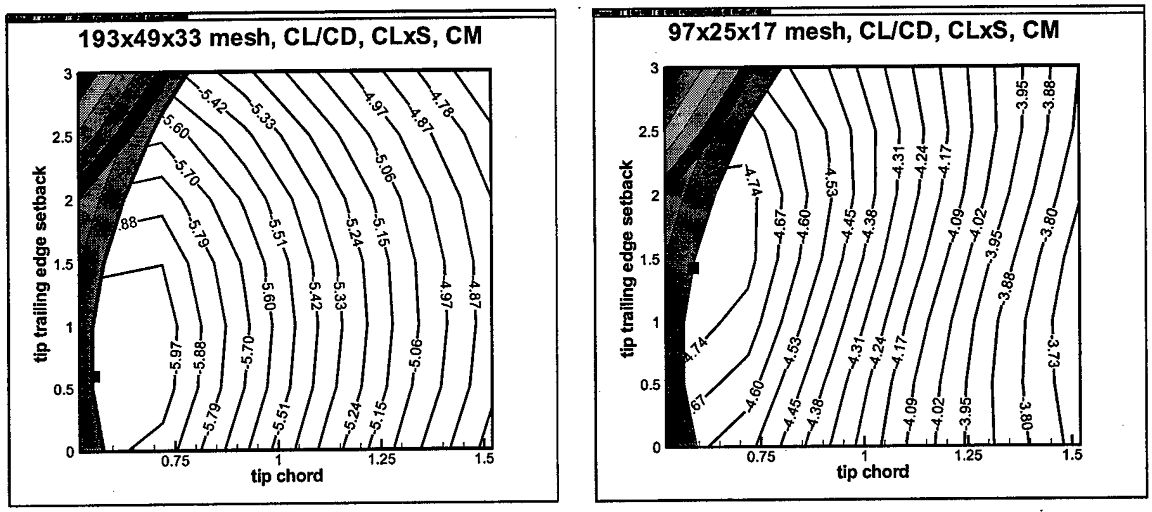

FIG. 4.2. *High-fidelity vs. low-fidelity objectives and active constraints: level sets of actual functions*

It was at that time that, despite a promising start, it was decided to postpone further studies with MAESTRO-based AMF because the test problem structure was inappropriate for the MAESTRO approach, given that the objective and constraint evaluations can be obtained only as a result of analysis, and there are few constraints and design variables. MAESTRO-based AMF will be tested at a later date on an bona fide MDO problem.

Computing with direct analyses was so laborious and time-consuming that at that stage it was also decided to build response surfaces out of the accumulated 193×49×33 and 97×25×17 data and to continue testing using these substitutes for the CFD analyses.

Figures 4.3, 4.4, and 4.5 show the resulting level sets of the response surface substitutes for the the objective and active constraints obtained from the same CFL3D.ADII data used to generate the level sets of the actual functions in Figure 4.2. While the objective functions are approximated well in all response surfaces, it is obvious that the spline and kriging approximations do well with constraints but a straightforward cubic polynomial response surface does not provide a good approximation to the problem constraints. Moreover, the low-fidelity polynomial model is not a good approximation to the high-fidelity polynomial model, as the Figure 4.5 demonstrates. Thus, the spline and kriging approximations model the situation in which the relationship between the high and low-fidelity approximations is favorable, while the cubic polynomial approximation models the situation when the relationship is not as favorable.

The augmented Lagrangian-based AMF was applied to a response surface substitute for the CFD analysis constructed via kriging. Conventional optimization required 37 evaluations of the high-fidelity objective and constraint values, and 27 evaluations of the high-fidelity objective and constraint sensitivities. The augmented Lagrangian AMF required 6 evaluations of the high-fidelity objective and constraint values, 6 evaluations of the high-fidelity objective and constraint sensitivities, 51 evaluations of the low-fidelity objective and constraint values, and 36 evaluations of the low-fidelity objective and constraint sensitivities. Since the low-fidelity analyses take 1/8 of the time of the high-fidelity analyses, the augmented Lagrangian required

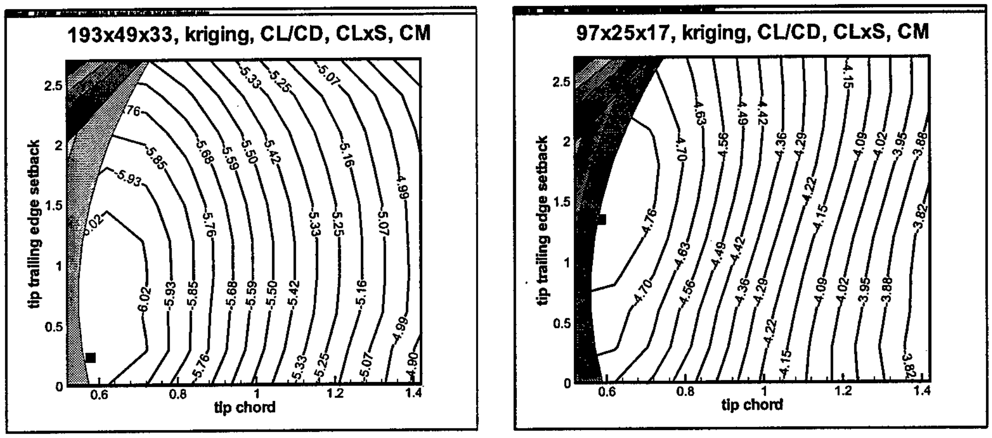

FIG. 4.3. *High-fidelity vs. low-fidelity objectives and active constraints: level sets of kriging approximation*

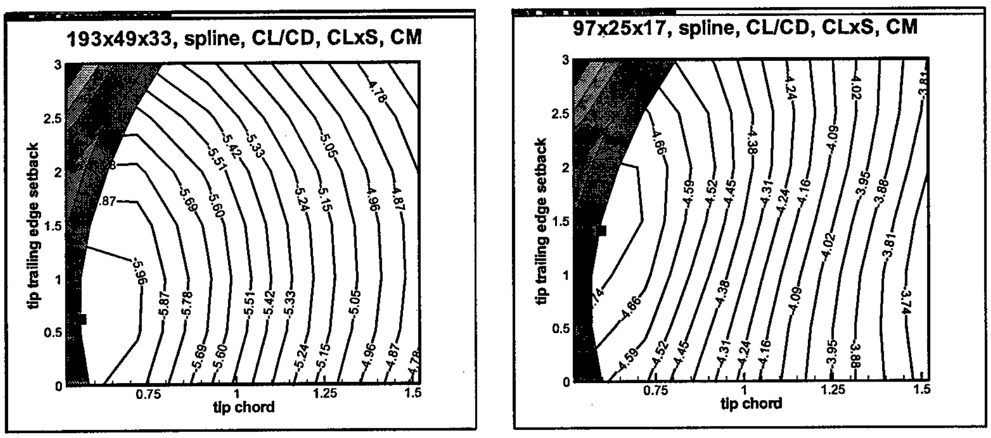

FIG. 4.4. *High-fidelity vs. low-fidelity objectives and active constraints: level sets of spline approximation*

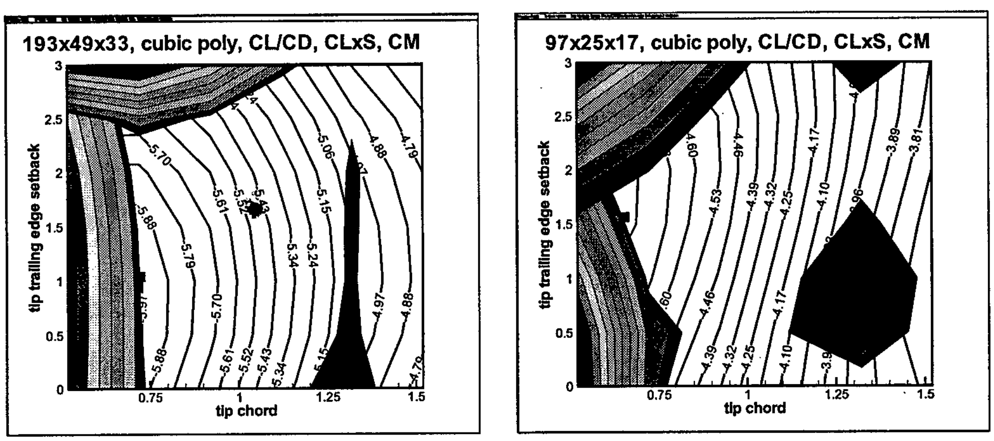

FIG. 4.5. *High-fidelity vs. low-fidelity objectives and active constraints: level sets of cubic polynomial approximation*

the equivalent work of $6 + 51/8 = 12\ 3/8$ evaluations of the high-fidelity objective and constraint values, and

$6 + 36/8 = 10\ 1/2$ evaluations of the high-fidelity objective and constraint sensitivities.

The SQP-based approach yielded similar improvements in performance. Conventional optimization, applied to a cubic polynomial substitute for the CFD analysis, required 31 high-fidelity functions and 31 high-fidelity sensitivities. Optimization using the SQP-based AMF required 4 high-fidelity functions and 51 low-fidelity functions, for a total of $4 + 51/8 = 10\ 3/8$ equivalent high-fidelity functions and as many sensitivities. For a spline substitute for the CFD analysis, conventional optimization required 21 high-fidelity functions and as many sensitivities. The SQP-based AMF required 4 high-fidelity functions, 4 high-fidelity sensitivities, 28 low-fidelity analyses, and 28 low-fidelity sensitivities, or a total of $4+28/8 = 7\ 1/2$ equivalent high-fidelity functions and as many sensitivities.

Both the augmented Lagrangian-based AMF and the SQP-based AMF produced consistent improvements in efficiency compared to non-AMF versions of the same codes. Improvements in efficiency due to each AMF are summarized in Table 4.1. Furthermore, the performance of the SQP-based AMF can be improved by reducing the amount of optimization done using the low-fidelity approximation.

TABLE 4.1

Wing optimization problem: Summary of improvement factor due to the AMF in function (first number) and sensitivity (second number) computations.

	Full CFD analysis	Kriging	Spline	Polynomial
Augmented Lagrangian AMF		3.0 / 2.6		
SQP AMF			2.8 / 2.8	3.0 / 3.0
MAESTRO AMF	1.9 / 1.9			

5. 2D Airfoil Problem.

5.1. Optimization Problem. In this problem, the objective function is the negative lift-to-drag coefficient ratio, $-C_L/C_D$, and the single nonlinear constraint is that on the pitching moment coefficient C_M. Figure 5.1 depicts the two design variables, maximum camber and maximum thickness. The flow is transonic, with $M_\infty = 0.8$, and $\alpha = 0°$. Function and constraint values are obtained with the FLOMG code [36] (applied to the Euler equations) evaluated on a 129×33 mesh and a 257×65 mesh, with the former currently providing the lowest level of fidelity. Figure 5.2 depicts the level sets obtained directly from FLOMG on the 129×33 and 257×65 meshes, respectively. This problem also has structure favorable for AMF. While the optima are at different locations, both the low-fidelity functions exhibit the same general trends as do the high-fidelity functions. Figure 5.3 shows that the spline response surface provides an excellent approximation to the actual functions.

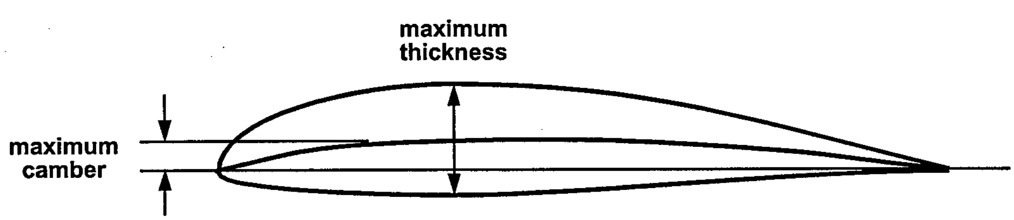

FIG. 5.1. *The 2D airfoil problem*

The time per analysis on the 257×65 mesh requires approximately four times the analysis time on the 129×33 mesh. On an SGI Octane workstation, the actual CPU times are approximately 8 and 2 minutes,

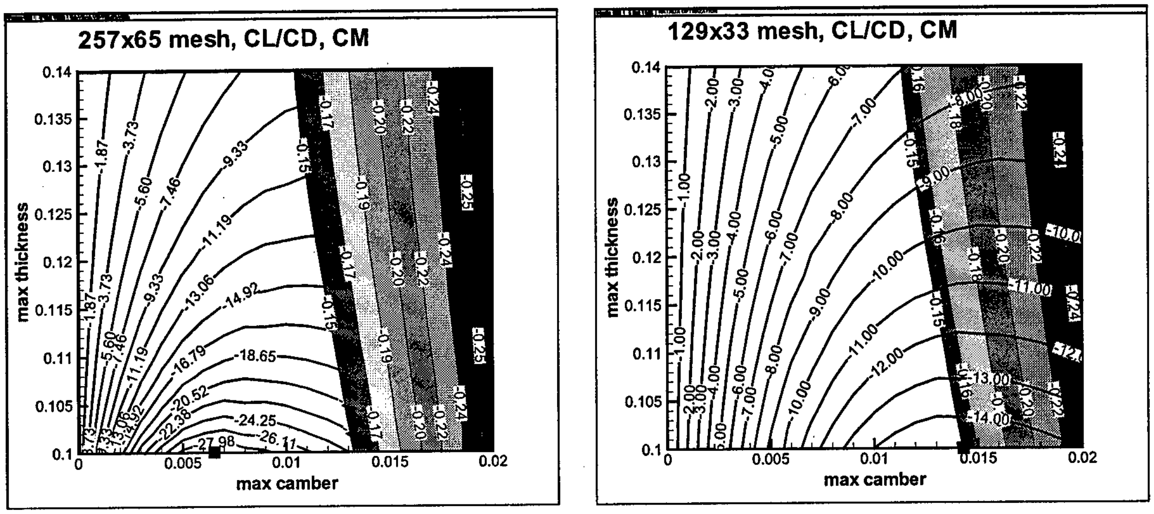

FIG. 5.2. *High-fidelity vs. low-fidelity objectives and active constraints: level sets of actual functions*

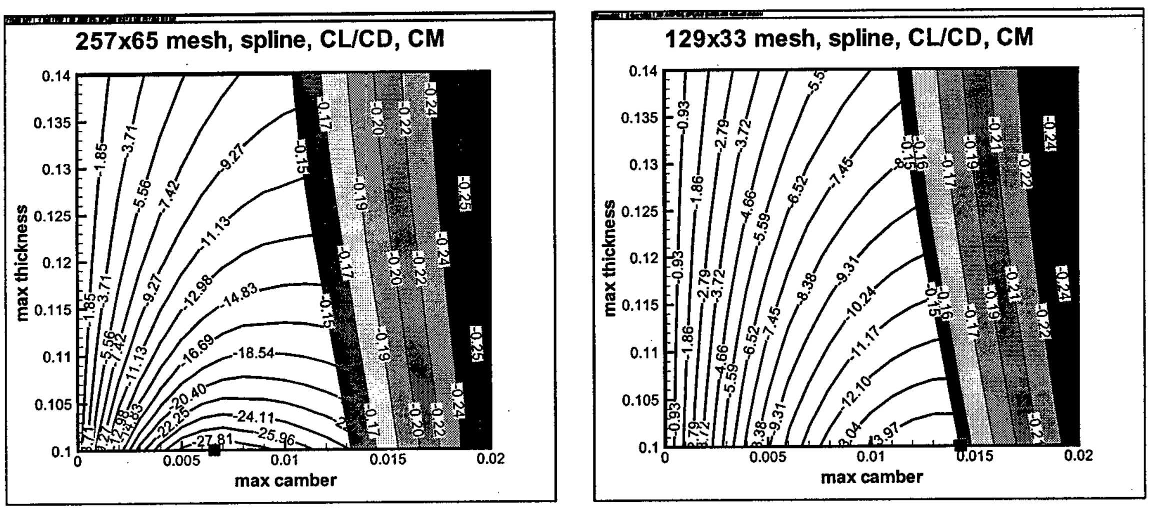

FIG. 5.3. *High-fidelity vs. low-fidelity objectives and active constraints: level sets of spline approximation*

respectively, iterating from free-stream conditions.

5.2. Discussion of Numerical Results. Again, the AMF consistently yielded improvements in efficiency compared to non-AMF versions of the same codes. However, since the airfoil problem is a 2D problem, the gains in relative efficiency are somewhat smaller (though still very good) than those observed for the 3D wing problem because the relative costs of the low- and high-fidelity calculations are smaller for the 2D calculations.

In tests done directly with FLOMG, MAESTRO required 34 iterations (each included an evaluation of the objective and constraints and their sensitivities) on the high-fidelity 257×65 mesh. The MAESTRO-based AMF required 20 iterations on the 129×33 mesh and 9 iterations on the 257×65 mesh. A reasonable comparison is made by considering that 20 iterations on the 129×33 mesh are equivalent to 5 iterations on the 257×65 mesh. Therefore, MAESTRO with AMF took 14 equivalent iterations.

The augmented Lagrangian-based AMF was applied to a spline substitute for the CFD analysis. Conventional optimization (using analytical derivatives) required 58 evaluations of the high-fidelity objective and constraint values, and 21 evaluations of the high-fidelity objective and constraint sensitivities. The augmented Lagrangian AMF required 6 evaluations of the high-fidelity objective and constraint values, 6 evaluations of the high-fidelity objective and constraint sensitivities, 50 evaluations of the low-fidelity objective and constraint values, and 30 evaluations of the low-fidelity objective and constraint sensitivities. Since the low-fidelity analyses take 1/4 of the time of the high-fidelity analyses, the augmented Lagrangian required the equivalent work of $6 + 50/4 = 18\ 1/2$ evaluations of the high-fidelity objective and constraint values, and $6 + 30/4 = 13\ 1/2$ evaluations of the high-fidelity objective and constraint sensitivities. These numbers yield approximately threefold improvement in the number of equivalent evaluations.

The SQP-based approach yielded similar improvements in performance. Conventional optimization, applied to the spline substitute for the CFD analysis required 19 high-fidelity function and sensitivity calls, each. Optimization using the SQP-based AMF required only 4 high-fidelity and 19 low-fidelity function and sensitivity calls, each, for a total of $4 + 19/4 = 8\ 3/4$ equivalent high-fidelity analyses. The 2D airfoil optimization results are summarized in Table 5.1.

TABLE 5.1

Airfoil optimization problem: Summary of improvement factor due to the AMF in function (first number) and sensitivity (second number) computations.

	Full CFD analysis	Spline
Augmented Lagrangian AMF		3.1 / 1.6
SQP AMF		2.2 / 2.2
MAESTRO AMF	2.4 / 2.4	

6. Concluding Remarks and Future Directions. In the preliminary experiments discussed here, the Approximation Management Frameworks (AMF) yielded about a threefold improvement in computational cost for the 3D wing design problem, and a twofold improvement for the 2D airfoil problem. It is believed that greater improvements can be achieved. No "fine tuning" of the AMF approaches has yet been done, and there is room for improvement in the interaction between all the pieces. In particular, currently the inner subproblem of minimizing the low-fidelity model is probably being solved to an unnecessarily high degree of accuracy. Because the analysis of the algorithms requires the subproblem solution to proceed only as far as needed to ensure sufficient improvement in the merit function of the high-fidelity problem, the subproblems are almost certainly being over-solved. The efficiency can thus be significantly improved if it is determined how to terminate the inner subproblem as soon as it produces the necessary decrease. This question is under investigation.

The efficacy of the AMF depends on the ability of the lower-fidelity model to predict the descent (or ascent) trends in the higher-fidelity model. It was found that even when this prediction was not favorable, as in the case of the cubic polynomial substitute, the first-order scaling technique due to Chang et al. provided an effective correction for the lower-fidelity model to ensure following the high-fidelity trends.

While these initial experiments are very promising, much work remains on further details of the implementation, as well as conclusions and practical guidance for using AMF for the selected modeling. The work includes study of the proper amount of optimization in the AMF subproblems and the consequences for overall efficiency of the interaction of the various levels of optimization in the AMF. The relative efficiency of AMFs based on different underlying optimization algorithms will also be studied. At this point the

SQP-based AMF is the most promising for the single discipline problems and for problem formulations that rely on multidisciplinary analysis. A variant of the augmented Lagrangian approach may have merit in the multidisciplinary setting as well. The MAESTRO approach is also promising for multidisciplinary problems. The AMF idea will also be applied to aerodynamic optimization with transonic flows; these problems should more fully exercise the AMF idea. Also to be examined are hierarchies of approximation based on models other than variable levels of discretization, such as direct response surface approximation of the high-fidelity model using kriging. Finally, the integration of the AMF idea in multidisciplinary problems will be studied.

Acknowledgments. The authors would like to thank Philip E. Gill for a copy of NPSOL and Anthony A. Giunta for a copy of RSG.

REFERENCES

[1] N. Alexandrov, J. E. Dennis, Jr., R. M. Lewis, and V. Torczon, *A trust region framework for managing the use of approximation models in optimization*, Structural Optimization, 15 (1998), pp. 16–23.

[2] N. M. Alexandrov, *Multilevel and multiobjective optimization in multidisciplinary design*, in Proceedings of the Sixth AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, September 1996. AIAA paper AIAA-96-4122.

[3] ——, *A trust region framework for managing approximation models in engineering optimization*, in Proceedings of the Sixth AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Optimization, September 1996. AIAA paper AIAA-96-4102.

[4] ——, *Multilevel methods for MDO*, in Multidisciplinary Design Optimization: State of the Art, N. M. Alexandrov and M. Y. Hussaini, eds., Society for Industrial and Applied Mathematics, Philadelphia, 1997, pp. 79–89.

[5] ——, *A trust-region framework for managing approximations in constrained optimization problems*, in Proceedings of the First ISSMO/NASA Internet Conference on Approximation and Fast Reanalysis Techniques in Engineering Optimization, June 14–27, 1998.

[6] N. M. Alexandrov and J. E. Dennis, Jr., *Multilevel algorithms for nonlinear optimization*, in Computational Methods for Optimal Design and Control, J. Borggaard, J. Burns, E. Cliff, and S. Schreck, eds., Birkhäuser, 1998, pp. 1–22.

[7] N. M. Alexandrov and R. M. Lewis, *First-order approximation management frameworks for nonlinear optimization.* To be published as an ICASE technical report.

[8] M. Avriel and A. C. Williams, *Complementary geometric programming*, SIAM Journal on Applied Mathematics, 19 (1970), pp. 125–141.

[9] J.-F. M. Barthelemy and R. T. Haftka, *Approximation concepts for optimum structural design—a review*, Structural Optimization, 5 (1993), pp. 129–144.

[10] C. Bischof, A. Carle, P. Khademi, and A. Mauer, *The ADIFOR 2.0 system for the automatic differentiation of Fortran 77 programs*, IEEE Computational Science and Engineering, 3 (1996), pp. 18–32.

[11] C. Bischof, W. T. Jones, J. Samareh-Abolhassani, and A. Mauer, *Experiences with the application of the ADIC automatic differentiation tool to the CSCMDO 3-D volume grid generation code.* AIAA paper 96-0716, 1996.

[12] F. G. Blottner and A. R. Lopez, *Determination of solution accuracy of numerical schemes as part*

of code and calculation verification, Tech. Rep. SAND98-222, Sandia National Laboratories, October 1998.

[13] A. J. BOOKER, J. E. DENNIS, JR., P. D. FRANK, D. B. SERAFINI, V. TORCZON, AND M. W. TROSSET, *A rigorous framework for optimization of expensive functions by surrogates*, Structural Optimization, 17 (1999), pp. 1–13.

[14] V. BRAIBANT AND C. FLEURY, *An approximation-concepts approach to shape optimal design*, Computer Methods in Applied Mechanics and Engineering, 53 (1985), pp. 119–148.

[15] S. BURGEE, A. A. GIUNTA, V. BALABANOV, B. GROSSMAN, W. H. MASON, R. NARDUCCI, R. T. HAFTKA, AND L. T. WATSON, *A coarse grained parallel variable-complexity multidisciplinary optimization paradigm*, The International Journal of Supercomputing Applications and High Performance Computing, 10 (1996), pp. 269–299.

[16] K. J. CHANG, R. T. HAFTKA, G. L. GILES, AND P.-J. KAO, *Sensitivity-based scaling for approximating structural response*, Journal of Aircraft, 30 (1993), pp. 283–288.

[17] A. R. CONN, N. I. M. GOULD, AND P. L. TOINT, *A globally convergent augmented Lagrangian algorithm for optimization with general constraints and simple bounds*, SIAM Journal on Numerical Analysis, 28 (1991), pp. 545–572.

[18] J. E. DENNIS, JR. AND R. B. SCHNABEL, *Numerical Methods for Unconstrained Optimization and Nonlinear Equations*, Prentice-Hall, Englewood Cliffs, NJ, 1983.

[19] P. E. GILL, W. MURRAY, M. A. SAUNDERS, AND M. H. WRIGHT, *User's guide for NPSOL (Version 5.0): A Fortran package for nonlinear programming*, Stanford University, 1995.

[20] P. E. GILL, W. MURRAY, AND M. H. WRIGHT, *Practical Optimization*, Academic Press, London, 1981.

[21] A. A. GIUNTA, *Aircraft Multidisciplinary Design Optimization Using Design of Experiments Theory and Response Surface Modeling Methods*, PhD thesis, Department of Aerospace and Ocean Engineering, Virginia Polytechnic Institute & State University, Blacksburg, Virginia, 1997; available as MAD 97–05–01, Multidisciplinary Analysis and Design Center for Advanced Vehicles, Virginia Polytechnic Institute & State University, Blacksburg, Virginia 24061–0203.

[22] P. HAJELA, *Geometric programming strategies in large-scale structural synthesis*, AIAA Journal, 24 (1986), pp. 1173–1178.

[23] A. JAMESON, *Essential elements of computational algorithms for aerodynamic analysis and design*, Tech. Rep. 97-68, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center, Hampton, Virginia, December 1997.

[24] L. KAUFMAN AND D. GAY, *PORT Library: Optimization and Mathematical Programming*, Bell Laboratories, May 1997.

[25] R. M. LEWIS, *A trust region framework for managing approximation models in engineering optimization*, in Proceedings of the Sixth AIAA/NASA/ISSMO Symposium on Multidisciplinary Analysis and Design, September 1996. AIAA paper 96–4101.

[26] J. J. MORÉ, *Recent developments in algorithms and software for trust region methods*, in Mathematical Programming. The State of the Art, Bonn 1982, A. Bachem, M. Grötschel, and B. Korte, eds., Springer-Verlag, Berlin, 1983, pp. 258–287.

[27] A. J. MORRIS, *Approximation and complementary geometric programming*, SIAM Journal on Applied Mathematics, 23 (1972), pp. 527–531.

[28] J. C. NEWMAN III, A. C. TAYLOR III, R. W. BARNWELL, P. A. NEWMAN, AND G. J.-W. HOU,

Overview of sensitivity analysis and shape optimization for complex aerodynamic configurations, Journal of Aircraft, 36 (1999), pp. 87–96.

[29] V. H. Nguyen, J. Strodiot, and C. Fleury, *A mathematical convergence analysis of the convex linearization method for engineering design optimization*, Engineering Optimization, 11 (1987), pp. 195–216.

[30] J. Rodríguez, J. E. Renaud, and L. T. Watson, *Trust region augmented Lagrangian methods for sequential response surface approximation and optimization*, in Proceedings of DETC '97, September 1997. ASME paper DETC97DAC3773, presented at the 1997 ASME Design Engineering Technical Conferences, September 14–17, Sacramento, California.

[31] C. Rumsey, R. Beidron, and J. Thomas, *CFL3D: Its history and some recent applications.* NASA TM 112861, 1997.

[32] L. A. Schmit, Jr. and B. Farshi, *Some approximation concepts for structural synthesis*, AIAA Journal, 12 (1974), pp. 692–699.

[33] L. A. Schmit, Jr. and C. Fleury, *Structural synthesis by combining approximation concepts and dual methods*, AIAA Journal, 18 (1980), pp. 1252–1260.

[34] L. A. Schmit, Jr. and H. Miura, *Approximation concepts for efficient structural synthesis.* NASA CR–2552, March 1976.

[35] R. E. Smith, M. I. G. Bloor, M. J. Wilson, and A. T. Thomas, *Rapid airplane parametric input design (RAPID).* AIAA paper 95-1687, 1995.

[36] R. C. Swanson and E. Turkel., *A multistage time-stepping scheme for the Navier-Stokes equations.* AIAA paper 85-0035, 1985.

[37] A. C. Taylor III, A. Oloso, and J. C. Newman III, *CFL3D.ADII (version 2.0): An efficient, accurate, general-purpose code for flow shape-sensitivity analysis.* AIAA paper 97-2204, 1997.

[38] Virginia Polytechnic Institute and State University, *RSG: Response Surface Generation Program, Version 1.32*, January 1996.